WORLD METEOROLOGICAL ORGANIZATION

COMMISSION FOR BASIC SYSTEMS
OPEN PROGRAMME AREA GROUP
ON INTEGRATED OBSERVING SYSTEMS

WORKSHOP ON RADAR DATA EXCHANGE

EXETER, UK, 24-26 APRIL 2013

CBS/OPAG-IOS/WxR_EXCHANGE/5.2

25.III.2013

ITEM: 5

Original: ENGLISH

CONSTRAINTS TO INCREASED WEATHER RADAR DATA EXCHANGE

Practical experiences of other global and regional observing system operators in overcoming data exchange constraints

(*R Giraud, ECMWF and Secretariat*)

SUMMARY AND PURPOSE OF DOCUMENT

Examples of approaches to data exchanges used by other communities.

ACTION PROPOSED

The meeting is invited to take the experiences of other data users into account.

Practical experiences of other global and regional observing system operators in overcoming data exchange constraints

1. Introduction

Advances both in technology and in science are driving an increase in the amount of information that is being requested from operators of weather radar systems. Added to this is the need to combine data from radars across continents to produce consistent composite coverages. This explosion in the amount of information to exchanged resembles the growth in the volume of information from satellites.

This paper introduces the main telecommunications technologies, together with a summary of their key attributes, and also describes how other science communities handle their need to exchange information.

2. Telecommunications technologies

Key telecommunications transmission methods, are outlined in Table 1.

Table 1. Summary of telecommunications technologies used for bulk data exchange

Technology	Description	Key strengths and weaknesses	Cost
Mobile telephony	Send and receive data over the mobile telephone network.	• Can be used where cabled access unavailable • Carried on internet • Availability affected by other users • Normally expensive	Low start-up Data volume charge high
Satellite broadcast	Usually uses "digital video broadcast" or "digital audio broadcast" standards. Files or streams of information are uplinked to a satellite and broadcast over a broad area.	• Typically used to broadcast information to a wide area • Uses "consumer technology" for reception • Economics depend on number of users receiving the same data	High start-up Cost of equipment only for extra users
Satellite upload	Similar to mobile telephony for transmitting information, but uses uplink to satellite and so has almost global coverage. Depending on the technology, dedicated fixed transmitting equipment may be required. Iridium, though, operates more like a terrestrial mobile telephone service.	• Limited transmission rates • Expensive for large volumes of data	Modest start-up Data volume charge high
Internet	The public internet	• There is no service level commitment for the core of the internet • Pricing and resilience of local connections depend on the chosen Internet Service Supplier. • The core internet itself is resilient. • There is no contact point to	Low start-up (but in some locations may need to lay cables). Charging policy differs between countries.

Technology	Description	Key strengths and weaknesses	Cost
		resolve problems. • Open to everyone; "denial of service" attacks can disrupt service • Technical tools allow prevention of data being tampered with *en route*. • Few protections against "denial of service" attacks	
Research network	These usually provide two sets of functionality: a) the research network operator acts as the user's Internet Service Provider b) features for handling very large data volumes (such as high bandwidth lines shared by only a small number of users)	• Usually higher bandwidth for a given cost than the public internet. • Core networks generally high availability (eg GEANT), but connections to the core depend on national research network arrangements. • Costs, and permission to connect, depend on national cost recovery policies. • No contact to manage problems with the core network (and sometimes with local connection) • High (many Gbit/s) bandwidths possible	Start-up depends on local infrastructure (laying cables can be expensive). Charging policy differs between countries.
Managed cloud – shared	A commercial network offered by an organization that acts like a private internet. Many customers may be active on the network, but each has a service level agreed and the supplier is obliged by the contract to delivery the network to the contract standards (this might include minimum bandwidth, maximum down time, etc). Typically, a supplier would define a community served by the cloud and only allow customers within that community to connect with each other through the cloud. In practice, many "fixed lines" are now delivered between fixed endpoints across a shared managed cloud.	• More predictable capacity that the internet • Single operator responsible for identifying and resolving faults • Shared with other user communities • Usually More expensive that internet solution • Supplier, rather than customer, decides on membership of the cloud. • Because the cloud is a "closed" community, and the traffic is managed by the supplier, the risk of denial of service attacks is lower than for the internet.	Low start-up Charging policy depends on country and level of service. More expensive than equivalent internet.
Managed cloud - private	This is like a shared managed cloud, but the organizations sharing the cloud are closely related. An example is the	• Predictable capacity • Single point of responsibility for resolving problems • Capacity and availability	Low start-up (unless cables need to be laid)

Technology	Description	Key strengths and weaknesses	Cost
	RMDCN. Although the underpinning technology infrastructure is shared among many user communities, the data flows of those communities are separated by a number of techniques so that the supplier can offer availability and bandwidth guarantees.	guarantees allow service planning • Usually more expensive that preceding options	Charging policy depends on country and level of service. More expensive than equivalent internet.
Managed private circuit	Single point-to-point link between two sites. Although sold as a private link it will run over shared infrastructure	• Single point of responsibility for resolving problems • Link can be "diverse routed" to reduce risk of line failure • Expensive • Each pair of centres needs to negotiate its membership • increasingly difficult to negotiate international links.	Low start up (unless cables need to be laid) Similar to public cloud (suppliers now use same technologies)
Dark Fibre	Use existing telecommunications infrastructure (normally fibre optic cables run alongside a major road, railway or powerline). The supplier provides minimal service other than connection to the fibre from the customer's site. Fault resolution normally falls to the customer	• Costs generally lower than other techniques • Fault resolution can be difficult • User normally has to design and negotiate resilience, often with different suppliers.	Start-up can be high because may need to run cables to access point. Ongoing charges low; may be local data charges if use local telecoms provider to connect to the fibre access point.

3. Exchange strategies

Although there are many strategies for how processing and exchange of observation information takes place, Table 2 outlines the key approaches on which other strategies are based.

Table 2. Strategies for computation and data exchange

Strategy	Description	Key strengths and weaknesses
Process at producer	All processing is performed on the same site as the observations are made. Only the products are exchanged. Example: automatic weather station preparing a bulletin containing a METAR and sending it to AFTN.	• Only products are transmitted (these tend to be smaller than the original data) • Need to operate and maintain computer systems and program suites at the observing sites • Limited ability for users to add further value to the products

Strategy	Description	Key strengths and weaknesses
		• Distributing updates to the processing software or exchange standards is difficult • Archive of input data difficult to access and manage (often not retained) • not feasible to reprocess data if recalibration needed • Change of observing technology may change products for users.
Process regionally	Minimal processing performed at the observing site. Data transferred to a regional centre for processing and products distributed from that centre. Example: Accounts for a group of companies; each company produces its own statement of accounts and these are combined to form the group statement of accounts	• Uniform processing within regions • Archiving managed at regional level • Data transfer within region greater than for local processing, but still low outside region • Limited ability to add value to products shared between regions • Can reprocess information from observing site if recalibration is needed • Change of technology at observing sites does not change products for users
Process centrally	Data are transferred to central point for processing and archive. Example: UK automatic weather station network.	• Reduced complexity and computers at remote sites, so easier to maintain and update • Increased data flow to centre • Changes to processing applied to all sites simultaneously • Changes to products made consistently • Uniform archive • Added value products possible because underlying information available to create derived products • Can reprocess data from archive if errors found in software or if calibration of observing equipment reviewed • Change of technology at observing sites does not change products for users • large data flow to centre; more limited data flow to end users.
Process at consumer	All data are transferred to the end user with minimal processing. Example: stock market ticker feeds	• Greatest flexibility for end user • End user needs to understand data and run processing software

Strategy	Description	Key strengths and weaknesses
	(average prices with change over a short interval)	• Very difficult to change observing systems because users have built systems around data stream characteristics • large data flow to end users.

3. Approaches used by science communities

3.1. Radio Astronomy – Very Long Baseline Interferometry (VLBI)

(http://www.evlbi.org/)
Very Long Baseline Interferometry uses a network of radio telescopes to create observations that have higher resolution than is possible using a single radio telescope. The technique tags the signals recorded at each radio telescope with a precise time so that the data from several telescopes can be processed as if they came from a single telescope.

Originally, the data were archived to magnetic media and these were transferred to a central point for processing. Such off-line processing met most of the needs of the radio astronomy community but was not able to respond to events as they evolved.

The European VLBI Network was founded in 1980 by five European institutes and has since expanded. In 2004, the network introduced fibre-optic connections, and now uses the national research networks, interconnected by the pan-european research network GEANT (http://en.wikipedia.org/wiki/G%C3%89ANT). Each node uses a 1Gbit/s connection.

In the Wikipedia article for the consortium (http://en.wikipedia.org/wiki/European_VLBI_Network) it states that the challenges being addressed by the group are the "final mile" connection between the telescopes and the research networks, and using distributed computing rather than a central processor.

This is an example of the "centralised processing" strategy.

3.2 UK Lightning Detection

The architecture of the UK lightning detection system is similar to that of VLBI (http://www.wmo.int/pages/prog/www/IMOP/publications/IOM-94-TECO2006/2(7)_Nash_UK.pdf). The sensors are deployed across much of the globe and their observations time stamped. These are transmitted to a central processing location where they are correlated and products showing the locations of lightning strikes produced.

The UK lightning detection system differs from that of the astronomers in that the information is time critical, but of relatively low volume. The key concerns in the design of the telecommunications are thus: availability (sites have to be away from radio interference), latency (the information has to be processed very soon after the sensors detect a signal), and cost (especially as the sensors are located around the globe).

The UK system uses available technologies to connect the sensors to the centre. At many sensor sites, a standard ADSL connection is used, so that the data are exchanged using the internet.

Particular challenges in developing the network have been around arranging the telecommunications connection at the remote site; separate negotiations have been needed with the local telecommunications operator in each host country.

3.3 CERN

CERN (the European Centre for Nuclear Research) is based in Geneva, spanning the Swiss-French border (http://www.wmo.int/pages/publications/showcase/documents/Hemmer_adj.pdf). Its instruments take measurements at a rate of a petabyte per second. Although the users do not need a real time feed of data, it is impracticable to store data at that rate, so real time processing is needed. The first stage of processing takes place close to the instruments, and reduces the volume of information to a size that can be archived and distributed to the "tier 1" archive and computation centres. These are connected to CERN with 10 GByte/sec links using the research network. These Tier 1 centres then supply the data to Tier 2 centres where most of the research simulation and analysis takes place.

CERN has introduced a backup site for data archive and processing in Budapest. This is linked to CERN by a dual-routed 100GByte/sec connection. Dual routeing is needed to provide resilience. CERN's need for resilience does not arise from real time application of the data, but from the impracticability of storing and forwarding the information resulting from the huge data volumes involved. Although the electronics used to drive networks is highly reliable, the fibre optic cables are vulnerable to being dug up by construction work, so routeing the cables so that they are physically separated reduces the level of risk. Physical separation also reduces other risks associated with the environment in which network equipment is located, such as fire and theft.

3.4 EUMETSAT

Meteorological satellites, such as those operated by EUMETSAT, generate large volumes of information. Many user communities need access to information in near real time.

EUMETSAT uses a variety of data dissemination techniques (http://www.eumetsat.int/Home/Main/DataAccess/index.htm?l=en).

The nature of satellites is that only limited processing can be performed onboard the satellite, so instrument-level information is transmitted to earth for further processing.

Some users have a need for detailed processing very soon after the sensor records information, and these are catered for through direct broadcast of information from the satellite (which is how EUMETSAT receives the information itself). This is the "process at consumer" strategy. Many information users, though, use products at differing levels of processing that are performed by EUMETSAT (a central strategy) or its Satellite Application Facilities (a regional strategy). Many of these products are transmitted to a communications satellite and broadcast to users from there.

Although all the methods of dissemination used by EUMETSAT form part of the WMO Information System (WIS), very little of the EUMETSAT information is exchanged through the high-availability Global Telecommunications System. Most is distributed through the direct broadcasts or the internet.

3.5 ECMWF and the RMDCN

ECMWF does not make observations, but it does generate very large data volumes. In general, it distributes products rather than the raw model information, and so uses a "process at producer" strategy. Because of its need to distribute information to its members, ECMWF developed a private computer network that has evolved to become the Regional Meteorological Telecommunication Network (RMDCN). Through an agreement with WMO, the RMDCN also acts as the core network for the WMO Information System, and provides the backbone of meteorological communications within Europe.

The RMDCN is a "managed network" that allows centres on it to exchange information with any other centre. The core of the network is highly resilient, and the supplier guarantees availability and minimum bandwidths for each connection. Each centre can choose how resilient to make its local connection, and the service allows the internet to be used as a backup communications system.

At a technical level, the network could be used to exchange information at rates far greater than are actually implemented. The cost of using the network increases with the guaranteed bandwidth, and most centres have chosen to implement bandwidths that are much lower than the speed of their connection to the internet. Typically, RMDCN is used by meteorological services for exchanging operationally critical information, and the internet or other techniques for information that is less time critical.

4. Summary

There is a trade off between reducing the difficulty of telecommunications by processing information close to the sensor and the cost of computing (and particularly maintenance) if the processing is not performed centrally.

A managed service can provide very high levels of availability, and responsibility of resolving problems with the network lies with the network provider. However, the cost for a given bandwidth is higher than for a self-managed service.
The core of the internet is highly reliable, but to build it into a high availability system requires care in the design of the connection between the user site and the internet point of presence. Some suppliers provide a resilient connection to the internet.

Research networks typically offer higher bandwidths that commercial operators, and it is possible to use technologies to create "virtual private networks" across these. Access to the research networks has different restrictions and costs in different countries.

In all cases, radar installations are seldom located in highly populated areas with a good telecommunications infrastructure. A major challenge is, therefore, the "final mile" connection between the radar site and the telecommunications backbone to which it has to connect.

Technology plays only a small part in the decision on which technology to use to interconnect radar centres. the key aspects are: cost of installation of the chosen telecommunications; cost of bandwidth; cost of managing the network contract and operations.